

UNITED STATES PATENT APPLICATION

for

METHOD AND APPARATUS FOR FILTERING EMAIL

Inventor:

Jerome R. Bellegarda
Devang Naik
Kim E. A. Silverman

Prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN, L.L.P.
12400 Wilshire Boulevard
Seventh Floor
Los Angeles, California 90025
(503) 684-6200

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The present invention relates generally to message filtering. More particularly, this invention relates to email filtering using latent semantic analysis.

Copyright Notice/Permission

[0002] A portion of the disclosure of this patent document contains material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever. The following notice applies to the software and data as described below and in the drawings hereto: Copyright © 2000, Apple Computer, Inc., All Rights Reserved.

Background

[0003] As the use of computers and the Internet have proliferated, so too has the use of email. Many businesses and consumers use email as a prominent means of communication. Not surprisingly, the exponential growth of the medium has also attracted the interest of commercial email advertisers. Commercial email advertisers obtain email addresses from a variety of sources, for example, from email vendors, or from commercial web sites, often without the permission of the owners of the email addresses. The email addresses may then be used to promote the products and services of the commercial email advertisers, or of the parties they represent.

[0004] The result is a deluge of unsolicited email received by hapless email users. One method to deal with unsolicited email is for a user to manually select and delete the unsolicited email. Other methods provide for recognizing a message sent in bulk to multiple recipients, and to either discard or tag the message as a possible unsolicited message. Still other methods maintain a database of addresses of known senders of unsolicited email and on receipt of the email, automatically discard those received from the known senders of unsolicited email. Still other methods use key-word filters. This method provides for scanning the subject and/or the body of the email message for some pre-determined keywords, and if detected, the message may be either discarded or tagged as suspicious.

[0005] Despite the methods described above, commercial email advertisers use ingenious methods to frustrate the efforts of email recipients. For example, to defeat the detection of bulk email, the email messages may be routed through a maze of servers so that ultimately, the message does not appear to be a bulk emailing. To defeat the system that tracks the address of known senders of unsolicited messages, the originating address of the unsolicited email may be changed often. To confuse keyword filter methods, the subject field of the email may be deceitfully titled, for example, "In response to your query". Moreover, the key-word filtering method suffers from other significant problems, for example, when trying to filter out email messages from pornographic email advertisers using the word "sex", legitimate anatomical or biological articles that include the word "sex" may also be eliminated.

SUMMARY OF THE INVENTION

[0006] A method and apparatus for filtering messages, in particular email messages is described herein. According to one aspect of the present invention, the method comprises determining a first semantic anchor corresponding with a first group of messages, for example, legitimate email messages and a second semantic anchor corresponding with a second group of messages, for example, unsolicited email messages. The method further determines a vector corresponding with an incoming message, compares the vector with at least one of the first semantic anchor and the second semantic anchor to obtain at least one comparison value, and filters the incoming message based on the comparison value.

[0007] Embodiments of the invention may be represented as a software product stored on a machine-accessible medium (also referred to as a computer-accessible medium or a processor-accessible medium). According to one aspect of the invention, the machine-accessible medium includes instructions that, when executed by a machine causes the machine to perform operations comprising determining a first semantic anchor corresponding with a first group of messages, for example, legitimate email messages. The machine-accessible medium includes further instructions for determining a second semantic anchor corresponding with a second group of messages, for example, unsolicited email messages. The machine-accessible medium includes further instructions for determining a vector corresponding with an incoming message, compares the vector with at least one of the first semantic anchor and the second semantic anchor to obtain at least one comparison value, and filters the incoming message based on the comparison value.

[0008] According to one aspect of the invention, the invention may be represented as an apparatus, e.g. computer system. The computer system comprises a bus, a data storage device coupled to the bus and a processor coupled to the data storage device, said processor to perform a method that comprises determining a first semantic anchor corresponding to a first group of messages. The processor also determines a second semantic anchor corresponding to a second group of messages. The processor further determines a vector corresponding to an incoming message, compares the vector corresponding to the incoming message with at least one of the first semantic anchor and the second semantic anchor to obtain a first comparison value and a second comparison value. The processor filters the incoming message based on the first comparison value and the second comparison value.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

Figure 1 is a block diagram illustrating an email filtering system according to one embodiment of the present invention;

Figure 2 is a block diagram of the matrices and vectors used in finding semantic anchors;

Figure 3 is a flowchart illustrating a method used in filtering email according to one embodiment of the present invention;

Figure 4 illustrates a block diagram of a computing device for use with one embodiment the present invention.

Figure 5 illustrates a block diagram of one embodiment of the invention stored on a machine-accessible medium.

DETAILED DESCRIPTION

[0010] Described is a method and apparatus for filtering email using latent semantic analysis.

[0011] In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one of ordinary skill in the art that the present invention may be practiced without these specific details. In other instances, well-known architectures, steps, and techniques have not been shown to avoid unnecessarily obscuring the present invention.

[0012] Parts of the description may be presented using terminology commonly employed by those skilled in the art to convey the substance of their work to others skilled in the art. Also, parts of the description may be presented in terms of operations performed through the execution of programming instructions. As well understood by those skilled in the art, these operations often take the form of electrical, magnetic, or optical signals capable of being stored, transferred, combined, and otherwise manipulated through, for instance, electrical components.

[0013] The invention may utilize a distributed computing environment. In a distributed computing environment, program modules may be physically located in different local and remote memory storage devices. Execution of the program modules may occur locally in a stand-alone manner or remotely in a client/server manner. Examples of such distributed computing environments include local area networks, enterprise-wide computer networks, and the global Internet.

[0014] In addition, it should be understood that the programs, processes, method, etc.

described herein are not related or limited to any particular computer or apparatus nor are they related or limited to any particular communication network architecture.

Rather, various types of general purpose machines may be used with program modules constructed in accordance with the teachings described herein. Similarly, it may prove advantageous to construct a specialized apparatus to perform the method steps described herein by way of dedicated computer systems in a specific network architecture with hard-wired logic or programs stored in nonvolatile memory such as read only memory.

[0015] Various operations will be described as multiple discrete steps performed in turn in a manner that is helpful in understanding the present invention. However, the order of description should not be construed as to imply that these operations are necessarily performed in the order they are presented, or even order dependent. Lastly, repeated usage of the phrase "in one embodiment" does not necessarily refer to the same embodiment, although it may.

[0016] Latent semantic analysis (LSA) is a method that automatically uncovers the salient semantic relationships between words and documents in a given corpus. Discrete words are mapped onto a continuous semantic vector space, in which clustering techniques may be applied. The method for filtering email messages comprises determining a first semantic anchor corresponding with a first group of email messages, for example, legitimate email messages and a second semantic anchor corresponding with a second group of email messages, for example, unsolicited email messages. Determining a vector corresponding with an incoming

email message, comparing the vector with at least one of the first semantic anchor and the second semantic anchor to obtain at least one comparison value, and filtering messages based on the comparison value.

[0017] **Figure 1** is a block diagram illustrating an email filtering system **100** according to one embodiment of the present invention. Although the description that follows describes the filtering of email messages, one skilled in the art will appreciate that the system may be used to filter email attachments, transcribed audio messages, computer programs, e.g., computer viruses, text, and the like. In one embodiment, email filtering system **100** filters unsolicited email messages from legitimate email messages. However, one skilled in the art will appreciate that other embodiments may classify messages into more than two groups.

[0018] Email filtering system **100** comprises an email training unit **105** that includes an email training corpus \mathcal{T} , for example, a database comprising a collection of N_1 legitimate email messages and N_2 unsolicited email messages. In one embodiment, the legitimate and unsolicited email messages are obtained from the existing email received by a recipient. Alternate embodiments may allow for a user to manually classify each incoming email message until an adequate email training corpus \mathcal{T} has been established. The words used in the collection of the legitimate email messages, and in the collection of the unsolicited email messages are from some underlying vocabulary v comprising, for example, the M most frequently used words in the language. In one embodiment, M may be ten thousand, and $1 \leq N_1 = N_2 \leq 150$.

[0019] Co-occurrences unit 110 of email filtering system 100, comprises a two dimensional ($M \times 2$) matrix W formed using the email training corpus \mathcal{T} . Matrix W essentially keeps track of which word is found in what document by keeping a record of the number of times each word appears in each legitimate and each unsolicited email message. In particular, entries ω_{ij} of matrix W reflects the extent to which each word $\omega_i \in \mathcal{V}$ appeared in the legitimate email message ($j=1$), or in an unsolicited email message ($j=2$). Various methods may be used to keep a record of the number of occurrences of a word in a document, for example, a simple normalized count of the number of occurrences of each word. However, in one embodiment, co-occurrence unit 110 uses function

$$\omega_{i,j} = (1 - \varepsilon_i) \frac{c_{i,j}}{N_j} \quad (1)$$

that normalizes for document length and word entropy to form matrix W . $c_{i,j}$ denotes the number of times each word ω_i occurs in the collection of legitimate email messages, and the number of times each word ω_i occurs in the collection of unsolicited email messages. In equation (1) N_j , for $j = 1$ and $j = 2$, represents the total number of words in the collection of legitimate email messages, and unsolicited email messages. ε_i is the normalized entropy of ω_i in the training email corpus \mathcal{T} . $(1 - \varepsilon_i)$ is merely a weighting factor, or a word distribution factor, and is a measure of the distribution of a particular word in the email training corpus \mathcal{T} . This is explained further below.

[0020] In one embodiment, co-occurrences unit 110 calculates ε_i using equation:

$$\varepsilon_i = -\frac{1}{\log N} \sum_{j=1}^N \frac{c_{i,j}}{t_i} \log \frac{c_{i,j}}{t_i} \quad (2)$$

where $N = N_1 + N_2$. By definition, $0 \leq \varepsilon_i \leq 1$, with equality if and only if $c_{i,j} = t_i$ and

$c_{i,j} = \frac{t_i}{N}$. Thus, a value of ε_i close to 1 indicates a word distributed across many email

messages throughout the email training corpus \mathcal{T} . However, a value of ε_i close to 0

indicates that the word is present only in a few email messages. Thus, the weighting

factor is a measure of the distribution of a word across the training email corpus \mathcal{T} . In

particular, weighting factor $(1 - \varepsilon_i)$ is a measure of the indexing power of the word ω_i .

[0021] After co-occurrences unit **110** constructs matrix W , Singular Value

Decomposition (SVD) unit **115** decomposes matrix W , and subsequently obtains the

semantic anchors \bar{v}_1 and \bar{v}_2 . The semantic anchors \bar{v}_1 **120** and \bar{v}_2 **125** are vectors

derived from matrix W using SVD. In one embodiment, vectors \bar{v}_1 and \bar{v}_2 are derived

using the following equation:

$$W = USV^T \quad (3)$$

where U is the $(M \times 2)$ left singular matrix with row vectors u_i ($1 \leq i \leq M$), S is the (2×2)

diagonal matrix of singular values $s_1 \geq s_2 > 0$, V is the (2×2) right singular matrix with

row vectors v_j ($j = 1, 2$), and T denotes matrix transposition. Thus, vector \bar{v}_1 represents

legitimate email messages and vector \bar{v}_2 represents unsolicited email messages.

[0022] **Figure 2** is a block diagram of the SVD of matrix W . As **Figure 2** illustrates, the

SVD of matrix W defines a mapping between the mathematical representation of

the set of legitimate and unsolicited email messages **205** and **210** respectively, and the latent semantic vector space spanned by the singular vectors contained in U and V.

The mapping is then scaled by the diagonal matrix **230** , to ensure proper representation. From this mapping, the first semantic anchor given by

$$\bar{v}_1 = v_1 S \quad (4)$$

and the second semantic anchor given by

$$\bar{v}_2 = v_2 S \quad (5)$$

are obtained after appropriate scaling by the diagonal matrix S. One skilled in the art will appreciate that V_1^T **215** and V_2^T **220** are unscaled semantic anchors in (2x2) matrix V^T **235**, and may be easily converted to 2-dimensional vectors \bar{v}_1 and \bar{v}_2 using the equations (4) and (5) above. If more than two classification groups are desired, i.e., classification groups other than legitimate and unsolicited, one skilled in the art will appreciate that semantic anchors corresponding to each classification group may be obtained as described above. Matrix U **240** is used to calculate the vector corresponding to an incoming email message as explained below.

[0023] Returning to **Figure 1**, whenever an incoming email message is received by incoming email unit **150**, equation 1 may be used by incoming email conversion unit **155** to convert the incoming email to a column vector d_3 of dimension M. In one embodiment, the resulting column vector d_3 may be inserted as an additional column in matrix W, thereby converting matrix W of dimension (Mx2) into a matrix of dimension (Mx3). Using the SVD of equation (2), an unscaled representation v_3^T of the new email

message is obtained. Thus, $d_3 = USv_3^T$, and hence the vector representation of the incoming email is obtained as follows:

$$\bar{v}_3 = v_3 S = d_3^T U \quad (6)$$

[0024] The 2-dimensional vector \bar{v}_3 of equation (6) is the mathematical representation of the new email message, and can be interpreted as a point in the latent semantic vector space spanned by vectors \bar{v}_1 and \bar{v}_2 .

[0025] One skilled in the art will appreciate that equation (6) is an approximate representation of the message in the existing LSA space. Since the new email message was not part of the original SVD extraction, words in the new email message, not found in training corpus \mathcal{T} , may cause the SVD expansion to no longer apply. As such, in one embodiment, an optional feed back path 180, as illustrated in **Figure 1**, may be used to add the new email message to the training corpus \mathcal{T} . Semantic anchors \bar{v}_1 and \bar{v}_2 may be periodically recalculated to account for the new words in the new email messages, so that subsequent email messages may be accurately classified as legitimate or unsolicited.

[0026] The invention contemplates capturing structural associations between words. Hence, two words whose representations are "close" (in some suitable metric) tend to appear in the same kind of documents, whether or not they actually occur within identical word contexts in the documents. Each semantic anchor \bar{v}_1 and \bar{v}_2 may be viewed as the centroid of the words in the legitimate email messages, and in the unsolicited email messages respectively. This means that associated words such as synonyms occur in close proximity to other similar words in each category of the unsolicited and legitimate

email messages in the semantic vector space S. For example, if a particular word is found more frequently in the unsolicited email messages as compared with legitimate email messages of the training corpus, an incoming email containing a synonym of the word will be closer to the unsolicited message category in semantic vector space S. Thus, email filtering system 100 properly classifies incoming email messages containing synonyms eliminating the need for frequent recalculations of semantic anchors \bar{v}_1 and \bar{v}_2 .

[0027] After calculating the semantic anchors \bar{v}_1 , \bar{v}_2 , and the vector representation \bar{v}_3 of the new email message, a measure of closeness K is calculated. The measure of closeness K is a measure of how close a new email message is to a legitimate email message, or to an unsolicited email message. The measure of closeness K is computed by calculation unit 160 and, in one embodiment, compares the angle formed between vectors \bar{v}_1 and \bar{v}_3 , with the angle formed between vectors \bar{v}_2 and \bar{v}_3 . The measure of closeness K may be calculated using:

$$K(\bar{v}_3, \bar{v}_j) = \cos(v_3 S, v_j S) = \frac{v_3 S^2 v_j^T}{\|v_3 S\| \|v_j S\|} \quad (7)$$

for $j = 1, 2$. Other methods may be employed to calculate the measure of closeness K including, but not limited to, calculating the length of the normals between vectors \bar{v}_1 , \bar{v}_2 , and \bar{v}_3 .

[0028] After calculating the measure of closeness K, logic unit 165 determines whether the new email is unsolicited 170, legitimate 175, or ambiguous 180. In one embodiment, if \bar{v}_3 is closer to \bar{v}_1 , i.e., the angle between \bar{v}_3 and \bar{v}_1 is smaller than the angle between \bar{v}_3 and \bar{v}_2 , the new email is considered to be a legitimate email

message 175, and email filtering system 100 may automatically permit the new email to be viewed by its intended recipient. Optionally, the email filtering system may allow the user to include the legitimate email message as part of the training email corpus \mathcal{T} . Alternately, if \bar{v}_3 is closer to \bar{v}_2 , i.e., the angle between \bar{v}_3 and \bar{v}_1 is greater than the angle between \bar{v}_3 and \bar{v}_2 , the new email is considered unsolicited 170. In one embodiment, unsolicited email messages may be automatically discarded by the email filter system. Alternate embodiments may maintain a copy of the unsolicited email so that a user may, at the user's convenience, discard the unsolicited mail or include it to form part of the training email corpus \mathcal{T} .

[0029] If the angle between \bar{v}_3 and \bar{v}_1 is approximately equal to the angle between \bar{v}_3 and \bar{v}_2 , logic unit 165 may tag the email message as ambiguous 180, for example, with an icon to indicate an ambiguous email message. Alternate embodiments may tag each incoming email message with a unique tag for each of the unsolicited, legitimate, and ambiguous categories, allowing for ease in sorting and handling of the received email messages. With respect to ambiguous email messages, in one embodiment, a user may determine whether the email message is legitimate or unsolicited. Alternate embodiments may permit a user to discard the ambiguous email message, or include it, after removing the ambiguity, to form part of the training email corpus \mathcal{T} , so that the ambiguity associated with future similar messages may be automatically handled by email filtering system 100.

[0030] As an example consider the following email messages received by a person in the fishing business: (a) Fishing is excellent in the south bank of the river, and (b) The Merchant bank has high interest rates. Although both email messages have the word 'bank' in the text of the message, the method described will properly classify message (a) as a legitimate email message and message (b) as an unsolicited email message.

[0031] The email training corpus \mathcal{T} is developed using existing email messages of the user in the fishing business. After the email training corpus \mathcal{T} is generated by email training unit 105, co-occurrences unit 110 generates matrix W using the email training corpus \mathcal{T} . SVD unit 115 decomposes matrix W and obtains semantic anchors \bar{v}_1 and \bar{v}_2 . When the two email messages are received by the user in the fishing business, they are each converted to a vector \bar{v}_3 using equation 6 above. In one embodiment, for each email message the measure of closeness K between \bar{v}_1 and \bar{v}_3 , and between \bar{v}_2 and \bar{v}_3 is calculated using equation 7. For legitimate email message (a), the measure of closeness K indicates that \bar{v}_3 is closer to \bar{v}_1 as compared with \bar{v}_2 thereby indicating the message is legitimate. However, for unsolicited message (b) the measure of closeness K will indicate that vector \bar{v}_3 is closer to unsolicited vector \bar{v}_2 as compared with \bar{v}_1 indicating that the message is unsolicited. Thus despite the same word 'bank' being present in each of the two email messages, the context in which they appear is taken into account in determining whether the received email message is legitimate or unsolicited.

[0032] Figure 3 illustrates a method that may be used to filter email according to one embodiment of the invention. At **305** the email training corpus \mathcal{T} is accessed, and at **310** the email messages in the training email corpus \mathcal{T} are used to construct matrix W (described earlier) that essentially keeps track of which word is found in what document. In particular, matrix W maintains a record of the number of times each word appears in each legitimate and each unsolicited email message. In one embodiment, equation (1) is used to construct matrix W . After constructing matrix W , at **315** a SVD is performed using equation (3) and semantic anchors \bar{v}_1 and \bar{v}_2 are obtained using equations (4) and (5).

[0033] At **320**, an incoming email message is received, and at **325**, vector \bar{v}_3 is constructed from the incoming email message using equation (6). At **330**, a measure of closeness K is obtained using equation (7). As explained above, the measure of closeness determines whether the new email message is legitimate, unsolicited or ambiguous.

[0034] At **335**, a determination is made whether the new email message is legitimate. If the angle between \bar{v}_3 and \bar{v}_1 is smaller than the angle between \bar{v}_3 and \bar{v}_2 , at **345**, the new email message is classified as legitimate. In one embodiment, legitimate email messages may be forwarded to the intended recipient.

[0035] At **350** a determination is made whether the new email message is unsolicited. If the angle between \bar{v}_3 and \bar{v}_1 is larger than the angle between \bar{v}_3 and \bar{v}_2 , at **340**, the new email message is classified as unsolicited. In one embodiment, the new email

message that is classified as unsolicited may be automatically discarded. Alternate embodiments may provide for the newly classified legitimate and unsolicited messages to form part of the email training corpus \mathcal{T} .

[0036] However, if the angle between \bar{v}_3 and \bar{v}_1 is approximately equal to the angle between \bar{v}_3 and \bar{v}_2 , at 355 the email message may be classified as ambiguous. In one embodiment, ambiguous email messages are forwarded to the intended recipient of the email message to eliminate the ambiguity and to classify the email message as legitimate or unsolicited. In one embodiment, after a recipient classifies the email message, the email message is included in the email training corpus, and new semantic anchors are calculated. Thus, the next time an email message with content similar to the ambiguous email message is received, the email filtering system automatically classifies the email as legitimate or unsolicited.

[0037] Embodiments of the email filtering system may be employed individually on a machine for a particular user or on a central machine, e.g., an email server, to filter out email messages for a group of email recipients. Alternate embodiments may include employing the email filtering system on a server or other device that communicates with a remote user, for example, a user using a wireless device such as a wireless personal digital assistant (PDA) or wireless palm top computer, so that the limited memory of the wireless device is not unnecessarily filled with unsolicited email messages. Alternate embodiments may employ the email filtering system on the PDA and unsolicited messages may be discarded as soon as they are received.

[0038] **Figure 4** illustrates one embodiment of an apparatus that may be used to filter email messages. Although the embodiment described uses a personal computer, other devices including wireless devices such as cellular phones and personal digital assistants may also be used. One embodiment of the present invention may be implemented on a personal computer (PC) architecture. It will be apparent to those of ordinary skill in the art that alternative computer system architectures or other processor, programmable or electronic-based devices may also be employed.

[0039] In general, such computer systems as illustrated by **Figure 4** include a processor **402** coupled through a bus **401** to a random access memory (RAM) **403**, a read only memory (ROM) **404**, and a mass storage device **407**. Mass storage device **407** represents a persistent data storage device, such as a floppy disk drive, fixed disk drive (e.g., magnetic, optical, magneto-optical, or the like), or streaming tape drive. Processor **402** may be any of a wide variety of general purpose processors or microprocessors (such as the Pentium® processor manufactured by Intel® Corporation), a special purpose processor, or a specifically programmed logic device.

[0040] Display device **405** is coupled to processor **402** through bus **401** and provides graphical output for computer system **400**. Input devices **406** such as a keyboard or mouse are coupled to bus **401** for communicating information and command selections to processor **402**. Also coupled to processor **402** through bus **401** is an input/output interface **410** which can be used to control and transfer data to electronic devices (printers, other computers, etc.) connected to computer system **400**. Computer system **400** includes network devices **408** for connecting computer system **400** to a network **414** through which email messages may be received, e.g., from

remote device 412. Network devices 408, may include Ethernet devices, phone jacks and satellite links. It will be apparent to one of ordinary skill in the art that other network devices may also be utilized.

[0041] One embodiment of the invention may be stored entirely as a software product on mass storage 407. Another embodiment of the invention may be embedded in a hardware product, for example, in a printed circuit board, in a special purpose processor, or in a specifically programmed logic device communicatively coupled to bus 401. Still other embodiments of the invention may be implemented partially as a software product and partially as a hardware product.

[0042] Embodiments of the invention may be represented as a software product stored on a machine-accessible medium (also referred to as a computer-accessible medium or a processor-accessible medium) as illustrated in **Figure 5**. The machine-accessible medium may be any type of magnetic, optical, or electrical storage medium including a diskette, CD-ROM, memory device (volatile or non-volatile), or similar storage mechanism. The machine-accessible medium may contain various sets of instructions, code sequences, configuration information, or other data. Those of ordinary skill in the art will appreciate that other instructions and operations necessary to implement the described invention may also be stored on the machine-accessible medium. **Figure 5** illustrates a machine-accessible medium that includes instructions that when executed by a machine causes the machine to perform operations comprising determining a first semantic anchor 520 corresponding with a first group of messages, for example, legitimate email messages. Determining a second semantic anchor 525 corresponding with a second group of messages, for example, unsolicited

email messages. The first and the second semantic anchors are determined as described earlier using instructions that implement the email training unit **505**, instructions that implement the co-occurrences unit **510** and instructions that implement the singular value decomposition unit **515**. The machine-accessible medium includes further instructions for determining a vector corresponding with an incoming message and instructions for comparing the vector with at least one of the first semantic anchor **520** and the second semantic anchor **525** to obtain at least one comparison value. The vector corresponding with an incoming message is determined using instructions to implement the email conversion unit **555**. The instructions for comparing the vector **555** with at least one of the first semantic anchor **520** and the second semantic anchor **525** to obtain at least one comparison value comprise instructions that implement the calculation unit **560**. The machine-accessible medium includes further instructions to filter the incoming message based on the comparison value. The instructions to filter the incoming message based on the comparison value, comprises instructions for implementing logic unit **565**. In particular, the instructions to filter the incoming message comprises instructions to determine whether the incoming message is unsolicited email **570**, legitimate email **575**, or ambiguous email **580**.

[0043] Experiments conducted using one embodiment of the method and apparatus of the present invention revealed that for a database comprising one legitimate email message N_1 and one unsolicited email message N_2 in the training corpus \mathcal{T} the email filtering system performed reasonably well. An exponential increase in the

performance of the email filtering system occurred as the values of N_1 and N_2 approached 50. Subsequent increases in the values of N_1 and N_2 revealed a relative plateau in the performance of the email filtering system. In one embodiment, more than 95% of a user's incoming email messages were properly classified, with approximately less than 3% of the user's email messages being passed to the user for disambiguation. A significantly lower misclassification rate was observed as compared with the misclassification rate of prior art methods.

[0044] While there has been illustrated and described what are presently considered to be example embodiments of the present invention, it will be understood by those skilled in the art that various other modifications may be made, and equivalents may be substituted, without departing from the true scope of the invention. Additionally, many modifications may be made to adapt a particular situation to the teachings of the present invention without departing from the central inventive concept described herein. Therefore, it is intended that the present invention not be limited to the particular embodiments disclosed, but that the invention include all embodiments falling within the scope of the appended claims.